


PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 11/14	A1	(11) International Publication Number: WO 98/48347 (43) International Publication Date: 29 October 1998 (29.10.98)
(21) International Application Number: PCT/US98/07456 (22) International Filing Date: 15 April 1998 (15.04.98) (30) Priority Data: 08/844,480 18 April 1997 (18.04.97) US (71) Applicant: STORAGE TECHNOLOGY CORPORATION [US/US]; 2270 South 88th Street, Louisville, CO 80028-4309 (US). (72) Inventors: WHITE, Michael, Wayne; 1424 Cypress Circle, Lafayette, CO 80026 (US). TOMSULA, Patrick, James; 10834 West 79th Place, Arvada, CO 80005-3439 (US). (74) Agents: GRAZIANO, James, M. et al.; Duft, Graziano & Forest, P.C., Suite 140, 1790 - 30th Street, Boulder, CO 80301-1018 (US).		(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>

(54) Title: MVS DEVICE BACKUP SYSTEM FOR A DATA PROCESSOR USING A DATA STORAGE SUBSYSTEM SNAPSHOT COPY CAPABILITY**(57) Abstract**

The MVS device backup system functions to enable the data processor to manage the device backup function of a disk data storage subsystem in a manner that minimizes the expenditure of data processor resources. This is accomplished by the MVS device backup system determining the source device volume on the data storage subsystem, the target device volume on the data storage subsystem and identifying the extent of both. The MVS device backup system then transmits data to the data storage subsystem, representative of the assignment of DASD full tracks from the source device location on the data storage subsystem as well as DASD full tracks from the target (backup) device location on the data storage subsystem. The data processor based MVS device backup system then uses ECAM channel programs to instruct the data storage subsystem to perform the device backup operation using snapshot track pointer copy operations. Upon conclusion of the device backup operation by the data storage subsystem, the MVS device backup system updates the meta data required to complete the device backup operation.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

**MVS DEVICE BACKUP SYSTEM FOR A DATA PROCESSOR
USING A DATA STORAGE SUBSYSTEM
SNAPSHOT COPY CAPABILITY**

5

FIELD OF THE INVENTION

This invention relates to data storage subsystems, and, in particular, to an MVS device backup system which is resident on a data processor and which regulates the backup of an MVS device on to the virtual data storage devices of a data storage subsystem which is connected to the data processor. The MVS device backup system manages the designation of the source and target data storage volumes and activates the snapshot copy resources of the data storage subsystem to perform the device backup operation without the necessity of the data processor having to expend a significant amount of processing resources.

PROBLEM

It is a problem in the field of computer systems to efficiently create backup copies of the data sets of an MVS device, which backup copies represent a single point-in-time for all data resident on the source MVS device. In a typical computer system, data processors are connected to one or more data storage subsystems, which include disk drive memory systems. The data processors and their associated data storage subsystems therefore must manage the backup of data sets stored on the virtual volumes of the data storage subsystem in a manner that does not adversely impact the performance and efficiency of the data processor and which also ensures the integrity of the data.

The backup of device data in a traditional DASD data storage subsystem entails the data processor retrieving the device data from the data storage subsystem, then writing the retrieved device data to a designated backup data storage location in the data storage subsystem. In particular, as part of this process, device data backup utilities perform volume track level backup by executing a series of read and write channel programs. These channel programs read data from the data storage subsystem into memory on a data processor and then write the data back out to the backup data storage subsystem as a series channel programs from that data processor memory. This system is resource intensive, in that the data processor expends a significant amount of resources (CPU cycles) to achieve the device data backup operation. In addition, channel resources are needed to perform the data processor based device backup

operation. Furthermore, two significant difficulties with data storage subsystems are the time required to make backup copies of device data and the need to maintain the consistency of the device data during the time it takes to make backup copies of the device data. This is a significant issue when a backup copy must be
5 made of a large devices which may contain many devices or data bases which are the target of a single or common application programs.

Several further problems are encountered in the creation of backup copies. The creation of physical backup copies of device data requires the expenditure of data processor, data channel, data storage subsystem and backup data storage
10 device resources. In addition, the concurrent access to the device data may be hindered due to device request queuing. This is necessitated by the requirement that the backup data represents a single point-in-time copy of the data for the entire source device.

An alternative to the data processor controlled device copy operation
15 described above is the data storage subsystem snapshot copy operation described in U.S. Patent No. 5,410,667. This snapshot copy system creates a duplicate device pointer in a virtual track directory in the data storage subsystem to reference a device that a data processor has requested the data storage subsystem to copy. This enables the data processor to access the device via two virtual addresses
20 while only a single physical copy of the device resides in the data storage subsystem. The snapshot copy operation is effected without the involvement of the data processor, since it is managed solely within the data storage subsystem. However, the snapshot copy operation cannot be selected by the data processor, but is because it is solely within the purview of the data storage subsystem. In
25 addition, the snapshot copy operation is not available for use in creating a backup copy of an MVS device.

SOLUTION

The above described problems are solved and a technical advance achieved in the field by the MVS device backup system of the present invention which
30 functions to enable the data processor to manage the device backup function in a manner that minimizes the expenditure of data processor resources. This is accomplished by the MVS device backup system designating the source and

destination device volumes and then activating the snapshot copy resources of the data storage subsystem to perform the device backup operation without the necessity of the data processor being involved in the execution details of the operation. In addition, the instantaneous creation of a backup copy of the device data maintains a point-in-time image of the device data from the initiation of the device backup process until a physical backup copy of the device data is available.

The implementation of the MVS device backup system is data processor based, yet the DASD volume device backup is performed without using data processor CPU resources to perform the actual movement of the device data. Thus, the traditional data reads to data processor memory and the write channel programs are not utilized to copy the device data from a source device location to a target (backup) device location. Instead, the MVS device backup system determines the source device volume on the data storage subsystem, the target device volume on the data storage subsystem and identifies the extents of both. The MVS device backup system then transmits data to the data storage subsystem, representative of the assignment of DASD full tracks from the source device location on the data storage subsystem as well as DASD full tracks from the target (backup) device location on the data storage subsystem. The data processor based MVS device backup system then uses Extended Channel Access Method (ECAM) channel programs to instruct the data storage subsystem to perform the MVS device backup operation using the data storage subsystem snapshot track pointer copy operations. Upon conclusion of the device backup operation by the data storage subsystem, the MVS device backup system updates the meta data required to complete the device backup operation. Meta data is the supporting volume and device structures, stored in the data processor, that identify the devices and maintain the device status.

This MVS device backup system is a significant departure from the prior art operation of data storage subsystems since it does not require the expenditure of a significant amount of the data processor resources. In addition, the existing data storage subsystem snapshot copy capability is used to enable the data processor to control the copying of device data to designated backup data storage locations.

BRIEF DESCRIPTION OF THE DRAWING

Figure 1 illustrates in block diagram form the overall architecture of a data system which includes the MVS device backup system of the present invention; and

- 5 Figure 2 illustrates in flow diagram form the operational steps taken by the MVS device backup system.

DETAILED DESCRIPTION

- Figure 1 illustrates in block diagram form the overall architecture of a computer system 1 that incorporates the MVS device backup system BS of the present invention. The computer system 1 includes at least one data processor 11 to which are connected at least one data storage subsystem (DSS) 3, 3' that contains at least one and likely a plurality of data storage devices DS1-DS4 for reading and writing data onto data storage media for use by the data processor 11, 12. The data storage subsystem 3, 3' comprises a dynamically mapped virtual device data storage subsystem which implements a plurality of virtual data storage devices. In the preferred embodiment of the invention disclosed herein, the data storage subsystem 3 comprises a disk drive array data storage subsystem, although any data storage technology can be used to implement the data storage subsystem 3. For the purpose of simplicity of description, the disk drive array data storage subsystem example is used herein.
- 10 The data storage subsystem 3 dynamically maps between three abstract layers: virtual, logical and physical. The virtual layer functions as a conventional large form factor disk drive memory. The logical layer functions as an array of storage units that are grouped into a plurality of redundancy groups. The physical layer functions as a plurality of individual small form factor disk drives. The data storage subsystem 3 effectuates the dynamic mapping of data among these abstract layers and controls the allocation and management of the actual space on the physical devices. These data storage management functions are performed in a manner that renders the operation of the data storage subsystem 3 transparent to the data processor 11, which perceives only the virtual image of the data storage subsystem 3. A virtual device is therefore an entity addressable by data processor
- 15
- 20

Mapping Tables

- 25
- 30

11, with data processor-controlled content and data processor-managed space allocation. In this system, the virtual device consists of a mapping of a large form factor disk drive image onto a plurality of small form factor disk drives which constitute at least one redundancy group within the disk drive array. The virtual to
5 physical mapping is accomplished by the use of a Virtual Device Table (VDT) entry which represents the virtual device. The "realization" of the virtual device is the set of Virtual Track Directory (VTD) entries, associated with the VDT entry each of which VTD entries contains data indicative of the Virtual Track Instances, which are the physical storage locations in the disk drive array redundancy group that contain
10 the data records. The data storage management functions are performed in a manner that renders the operation of the data storage subsystems 3, transparent to data processors 11.

Device Snapshot Copy Operation

As described in U.S. Patent No. 5,410,667, the data storage subsystem 3
15 includes a device copy capability which is termed a "snapshot copy" operation. The device snapshot copy operation instantaneously creates a second instance of a selected device by merely generating a new pointer to reference the same physical memory location as the original reference pointer in the virtual track directory. In this fashion, by simply generating a new pointer referencing the same physical
20 memory space, the device can be copied. A physical copy of the original device can later be written as a background process to a second memory location, if so desired. Alternatively, when one of the programs that can access the device writes data to or modifies the device in any way, the modified track is written to a new physical memory location and the corresponding address pointers are changed to
25 reflect the new location of this rewritten portion of the device.

This apparatus therefore instantaneously copies the original device without the time penalty of having to copy the track to the data processor 11 and then write the track to a new physical memory location on the data storage subsystem 3 via data channel 21. For the purpose of enabling a program to simply access the track
30 at a different virtual address, the use of this mechanism provides a significant time advantage. In this fashion, a track can be instantaneously copied by simply creating a new memory pointer and the actual physical copying of the track can

take place as a background process without the involvement of the data processor 11.

Operation of MVS Device Backup System

Figure 2 illustrates in flow diagram form the operation of the MVS device backup system BS, as implemented in the system environment described above. The MVS device backup system BS is illustrated herein as a device management process using the snapshot track level subsystem copy facility which is extant on the data processor 11, although other implementations, such as merging the functionality of the MVS device backup system BS into the other software running on data processor 11, are possible. In operation, data processor 11 requests the creation of a backup copy of an MVS device at step 201, which MVS device comprises a plurality of volumes of data sets stored on a virtual device on a one (for example DSS 3) of the data storage subsystem connected to data processor 11. The MVS device backup system BS, in response to the receipt of a device backup request, which includes an identification of the selected (source) MVS device, locates the identified source MVS device at step 202 and serializes the device staging request. The MVS device backup system BS searches the DSCB entries of the VTOC available to the data processor 11 and executes device type queries where necessary to thereby determine the device characteristics of the identified source MVS device as designated by the data processor 11. However, the data storage subsystem 3 stores data for the data processor 11 in virtual volumes, a plurality of which are implemented in data storage subsystem 3, as for example 256 logical DASD volumes. Thus, the device location information obtained by the MVS device backup system BS comprises the data processor view of the data storage location in which the selected source MVS device is stored on data storage subsystem 3.

The MVS device backup system BS, in response to the receipt of a device backup request, also locates or allocates a target location for the storage of the copy of the selected source MVS device at step 203. The MVS device backup system BS searches the DSCB entries of the VTOC available to the data processor 11 and executes device type queries where necessary to thereby determine the physical location of available memory space of sufficient extent to store the copy

of the identified source MVS device. In addition, the selection of the online candidate backup device can be accomplished by user request or automatically selected based upon device usage criteria. The MVS device backup system BS typically queries the backup device to determine relevant information regarding the present state of this device, including device label, device characteristics, available memory, and the like. Once the backup device has been selected, the device is serialized by the MVS device backup system BS. The data storage subsystem 3 stores data for the data processor 11 in virtual volumes, a plurality of which are implemented in data storage subsystem 3 and the device location information obtained by the MVS device backup system BS comprises the data processor view of the data storage location in which the copy of the selected device will be stored on data storage subsystem 3, if data storage subsystem 3 is the target device for storage of the backup copy of the identified MVS device.

Once the source and target locations are identified, the MVS device backup system BS, at step 204, builds a set of ECAM channel programs to activate the copying of the identified source MVS device, and the associated creation of track pointers. In operation, the MVS device backup system BS creates an ECAM channel message for each extent of tracks that are to be copied in the identified source MVS device. These ECAM messages are submitted serially to the data storage subsystem 3 to effect the device backup operation. In particular, the MVS device backup system BS issues an ECAM source definition message to the data storage subsystem 3 to define both the Source and Target Locations as well as the Extent for both of these locations. The MVS device backup system BS transmits the ECAM source definition message to the data storage subsystem 3 over the data channel 21 which connects these two systems in well known fashion. The ECAM source definition message is used to prepare the data storage subsystem 3 for the execution of the device snapshot copy operation and also functions as a security measure. The security aspect of the process involves the passing of a token at step 205 from the data storage subsystem 3 to the MVS device backup system BS to tie the received data with future instructions received from the MVS device backup system BS. In response to the receipt of the ECAM source

definition message, the data storage subsystem 3 stores the received Source and Target Location data as well as the Extent data for future use.

The MVS device backup system BS at step 206 generates an ECAM Target Execution message which instructs the data storage subsystem 3 to execute the requested MVS device backup operation, using the Source and Target Location, and Extent data received in the previous ECAM source definition message. This message is tied to the previously transmitted ECAM source definition message by the inclusion of the token which was received from the data storage subsystem 3 in the previously executed ECAM source definition message exchange. The data storage subsystem 3 uses the token included in the ECAM Target Execution message to identify the stored Source and Target Location and Extent data to perform the MVS device backup operation in well known fashion. Since the data storage subsystem snapshot copy operation executes independent of the data processor, the data storage subsystem 3 returns an MVS device backup acknowledgment message, including a polling token, at step 207 to the MVS device backup system BS on data processor 11. This polling token enables the MVS device backup system BS to query the data storage subsystem 3 to ascertain the completion status of the MVS device backup process, which physical copy is created as a background process, as noted above.

The MVS device backup query is implemented via the use of an ECAM Status Inquiry message, which comprises a status inquiry message. The MVS device backup system BS inserts the polling token in the ECAM Status Inquiry message to identify the particular device snapshot copy operation of interest (which represents the MVS device backup) and transmits the message at step 208. The data storage subsystem 3 responds to a received ECAM Status Inquiry message by returning a response at step 209 indicative of snapshot copy (MVS device backup) completed status. The MVS device backup system BS can originate as many ECAM Status Inquiry messages as needed until the data storage subsystem 3 transmits a response to the ECAM Status Inquiry message which indicates successful completion of the device snapshot copy (MVS device backup) operation.

Once the MVS device backup operation is completed, the MVS device backup system BS at step 210 must update the meta data associated with the

Target backup MVS device. The meta data comprises: VTOC; Volume labels; VTOC index; VVDS data, which is data processor 11. The MVS device backup system BS updates this data in well-known fashion to accurately reflect the location and extent of both the source MVS device as well as the backup copy of the MVS device created by the device snapshot copy operation of the data storage subsystem 3. Once the MVS device backup system BS updates the meta data, the MVS device backup system BS forces device recognition processing to be activated for the backup MVS device, which thereby allows concurrent access to both the source MVS device and the backup MVS device. At this juncture, the MVS device backup operation is completed and processing exits at step 211.

Summary

The MVS device backup system therefore functions to enable the data processor to manage the device backup function of a disk data storage subsystem in a manner that minimizes the expenditure of data processor resources. This is accomplished by the MVS device backup system determining the source MVS device on the data storage subsystem, the target MVS device on the data storage subsystem and identifying the extents of both. The MVS device backup system then transmits data to the data storage subsystem, representative of the assignment of DASD full tracks from the source MVS device location on the data storage subsystem as well as DASD full tracks from the target MVS device location on the data storage subsystem. The data processor based MVS device backup system then uses ECAM channel programs to instruct the data storage subsystem to perform the MVS device backup operation using snapshot track pointer copy operations. This eliminates the need for the data processor to be intimately involved in the execution of the MVS device backup operation.

What is Claimed:

1. A device backup apparatus for the copying of of a selected data storage device which is stored on a dynamically mapped virtual memory data storage subsystem having a rewriteable memory space, which device backup apparatus is extant on a data processor connected to the data storage subsystem
5 which is operational to instantaneously create a copy of a selected device independent of said data processor, said device backup apparatus comprising:
 - means for allocating memory in said rewriteable memory space for a selected device written thereon as well as for a copy of said selected device;
 - means for transmitting control messages to said data storage subsystem to
10 identify all data sets which comprise said selected device, and memory space for said copy of said selected device;
 - means for activating said data storage subsystem to implement said instantaneous copy operation on said identified all said data sets which comprise said selected device; and
15 means, responsive to said data storage subsystem copying said selected device, for updating meta data associated with a copy of said selected device created by said data storage subsystem, which metadata is known by said data processor.
2. The device backup apparatus of claim 1 wherein said means for allocating memory comprises:
 - means for determining a location and an extent of said selected device.
3. The device backup apparatus of claim 1 wherein said means for allocating memory comprises:
 - means for determining a location and an extent of said copy of said selected device.
4. The device backup apparatus of claim 1 wherein said means for transmitting control messages comprises:

means for transmitting data to said data storage subsystem to identify a location and an extent of said selected device, and a location and an extent for said
5 copy of said selected device.

5. The device backup apparatus of claim 1 wherein said means for updating meta data comprises:

means for updating at least one of device meta data comprising: VTOC; Volume labels; VTOC index; VVDS data.

6. A device backup method for the copying of a selected data storage device which is stored on a dynamically mapped virtual memory data storage subsystem having a rewriteable memory space, which device backup method is extant on a data processor connected to the data storage subsystem which is
5 operational to instantaneously create a copy of a selected device independent of said data processor, said device backup method comprising the steps of:

allocating memory in said rewriteable memory space for a selected device written thereon as well as for a copy of said selected device;

transmitting control messages to said data storage subsystem to identify said
10 selected device, and memory space for said copy of said selected device;

activating said data storage subsystem to implement said instantaneous copy operation on said identified all data sets which comprise said selected device; and

updating, in response to said data storage subsystem copying said selected
15 device, meta data associated with said device and known by said data processor.

7. The device backup method of claim 6 wherein said step of allocating memory comprises:

determining a location and an extent of said selected device.

8. The device backup method of claim 6 wherein said step of allocating memory comprises:

determining a location and an extent of said copy of said selected device.

9. The device backup method of claim 6 wherein said step of transmitting control messages comprises:

transmitting data to said data storage subsystem to identify a location and an extent of said selected device, and a location and an extent for said copy of said

5 selected device.

10. The device backup method of claim 6 wherein said step of updating meta data comprises:

updating at least one of device data comprising: VTOC; Volume labels; VTOC index; VVDS data.





FIG. 1

2/2
FIG. 2



INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/07456

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F11/14

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 566 967 A (INTERNATIONAL BUSINESS MACHINES) 27 October 1993 see column 6, line 6 - line 54 ---	1-10
A	EP 0 767.431 A (STAC, INC.) 9 April 1997 see column 6, line 26 - line 52 -----	1-10

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

29 July 1998

Date of mailing of the international search report

06/08/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Corremans, G

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 98/07456

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 566967	A	27-10-1993	US 5448718 A	05-09-1995
			JP 2703479 B	26-01-1998
			JP 6083686 A	25-03-1994
<hr/>				
EP 767431	A	09-04-1997	JP 10055298 A	24-02-1998
<hr/>				